

EXHIBIT D

#15



PATENT
00-S-023

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:)
CHARLES F. NEUGEBAUER) Group Art Unit: 2672
Serial No.: 09/536,880)
Confirmation No.: 3367) Examiner: J. Brier
Filed: March 27, 2000)
For: CONTEXT SENSITIVE SCALING)
CIRCUIT AND METHOD)

)

RECEIVED

MAR 02 2005

Technology Center 2600

APPELLANT'S BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Appellant hereby respectfully submits his brief in support of his appeal to the Board of Patent Appeals and Interferences from the decision dated October 30, 2003 of the Examiner finally rejecting claims 3, 5-10, 12-16, 18-22, 24, and 25 of the above-referenced application.

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:
Commissioner for Patents, P.O. Box 1450,

Alexandria, VA 22313-1450, on 8/9/04
Date of Deposit

Stephen Boagini
Appellant, Assignee, or Representative

Signature

8/9/04
Dated

1. REAL PARTY IN INTEREST

The real party in interest is STMicroelectronics, Inc. of Carrollton, Texas.

2. RELATED APPEALS AND INTERFERENCES

There are no related appeals or interferences.

3. STATUS OF CLAIMS

Claims 2-25 are pending. Claims 3, 5-10, 12-16, 18-22, 24, and 25 were finally rejected in the Office Action dated October 28, 2003. Claims 2 and 4 were indicated as being allowable over the art of record, and claims 11, 17, and 23 were indicated as being allowable if rewritten to include all of the limitations of the base claim and any intervening claims. Claims 3, 5-10, 12-16, 18-22, 24, and 25 are on appeal.

Attached hereto is an Appendix containing a copy of claims 3, 5-10, 12-16, 18-22, 24, and 25, which are the claims involved in this appeal.

4. STATUS OF AMENDMENTS

Appellant has not filed any amendments subsequent to the final rejection in the Office Action dated October 30, 2003.

5. SUMMARY OF THE INVENTION

The present invention is directed to methods and devices for scaling an image from one resolution to another in which the convolution kernel to be used for the scaling is selected based on local image content. One embodiment of the present invention provides a method for scaling a source image to produce a destination image that has a different resolution than the source image, as shown in Figures 2 and 4. [page 6, line 22 through page 9, line 12] According to the method, a local context metric is calculated from a local portion of the source image S12. [page 6, lines 23-25; page 7, lines 4-14; page 8, lines 6-16] A convolution kernel is generated from a plurality of available convolution kernels based on the calculated local context metric S14, and the generated convolution kernel is used to generate at least one pixel of the scaled destination

image S16 [page 6, line 25 through page 7, line 2; page 8, lines 6-9 and 16-20; page 9, lines 5-12] Those available convolution kernels include at least one smoothing kernel and at least one sharpening kernel. [page 8, lines 6-9; page 9, lines 5-12; Fig. 4]

Similarly, another embodiment of the present invention provides an image scaling device that receives pixels of a source image and outputs pixels of a scaled destination image that has a different resolution than the source image, as shown in Figures 4 and 5. [page 9, line 5 through page 10, line 4] The image scaling device includes a context sensor 34, a kernel generator 36, and a scaler 32. The context sensor 34 calculates a local context metric based on local source image pixels. [page 9, lines 15-16; page 11, lines 8-13] The kernel generator 36 generates a convolution kernel from a plurality of available convolution kernels based on the calculated local context metric. [page 9, lines 5-12 and 16-20; Fig. 4] The scaler 32 receives the coefficients of the generated convolution kernel from the kernel generator 36, and uses those coefficients to generate at least one pixel of the scaled destination image from pixels of the source image. [page 9, lines 20-23; page 10, lines 12-24] Those available convolution kernels include at least one smoothing kernel and at least one sharpening kernel. [page 9, lines 5-12; Fig. 4]

Thus, both embodiments of the present invention generate a pixel in the scaled destination image using a convolution kernel that is selected from available convolution kernels that include a smoothing kernel and a sharpening kernel, based on the value of a local context metric. Because the pixels of the scaled image are generated by selectively sharpening or smoothing the source image depending on the local content, a higher quality scaled image is produced. For example, in one exemplary embodiment of the present invention a local contrast metric is used along with a gaussian smoothing kernel and a cubic sharpening kernel in order to selectively sharpen the text and smooth the graphics of a computer video image, so as to produce a higher quality scaled computer video image. [page 8, lines 6-22]

6. ISSUES

- A. Whether claims 3, 5, 6, 8, 10, 12, 14, 16, 18, 20, 22, and 24 are anticipated under 35 U.S.C. § 102(e) by Lin (U.S. Patent No. 6,044,178).
- B. Whether claims 7, 9, 13, 15, 19, 21, and 25 are unpatentable under 35 U.S.C. § 103(a) over Lin in view of Miyake (U.S. Patent No. 6,088,489).

7. GROUPING OF CLAIMS

Group I

Claims 3, 5, 6, 8-10, and 12 stand or fall together.

Group II

Claims 14-16, 18, 20-22, and 24 stand or fall together.

Group III

Claims 7, 13, 19, and 25 stand or fall together.

8. ARGUMENT

- A. CLAIMS 3, 5, 6, 8, 10, AND 12 ARE PATENTABLE OVER LIN AND CLAIM 9 IS PATENTABLE OVER LIN IN VIEW OF MIYAKE

Appellant respectfully submits that claims 3, 5, 6, 8, 10, and 12 (of Group I) are patentable over Lin and claim 9 (of Group I) is patentable over Lin in view of Miyake because neither Lin nor Miyake teaches or suggests a method for scaling a source image to produce a scaled destination image in which a pixel of the scaled destination image is generated by using a convolution kernel that is generated from a plurality of available convolution kernels that include at least one smoothing kernel and at least one sharpening kernel based on a calculated local context metric.

Claim 5 recites a method for scaling a source image to produce a scaled destination image in which a convolution kernel is generated from a plurality of available convolution kernels based on a calculated local context metric, and the generated convolution kernel is used to generate at least one pixel of the scaled destination image, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

The Lin reference discloses an image processing apparatus and method for translating a source image to a display image having lower resolution. In the image processing system of Lin, a source image is first separated into a black text image portion, a white text image portion, and a background image portion. This separation of the source image S is performed by the combination of a thresholding unit 52, a text masking unit 54, a text segmentation unit 58, and an image separation unit 66, as shown in Figure 2. The three separated image portions are then separately processed by a down-sampling unit 70 in order to produce corresponding portions of the down-sampled (or scaled) image, as shown in Figure 3.

More specifically, the down-sampling unit 70 includes a first filter unit 78 for scaling the black text portion of the source image, a second filter unit 86 for scaling the white text portion of the source image, and a third filter unit 72 for scaling the background image portion of the source image. The scaled black text image portion and the scaled white text image portion are then further processed by separate sigmoid filter units 82 and 90. An image merging unit 96 merges the three resulting scaled image portions to produce the complete scaled image to be displayed on an LCD panel 100.

While Lin does disclose using multiple filter units 78, 86, and 72 to separately scale different portions of the source image that each have a different type of image content, all filter units in the image processing system of Lin are low pass filters. In particular, Lin teaches that "low pass filter 78" and "low pass filter 86" each implement a "gaussian filter" with kernel coefficients of "1/15, 3/15, 7/15, 3/15, 1/15" for scaling the black text and white text image portions of the source image, and "low pass filter 72" implements "a gaussian filter" with kernel coefficients of "1/10, 1/5, 2/5, 1/5, 1/10" for scaling the background image portion of the source image. See Lin at 6:18-22, 6:33-37, 6:55-58.

In the context of image scaling, a "low pass filter" is a "smoothing filter" that uses a "smoothing convolution kernel" derived from a "smoothing function" to smooth (or blur) an image by removing detail and noise. In contrast, a "sharpening filter" uses a "sharpening convolution kernel" derived from a "sharpening function" to sharpen (or enhance the detail) of an image by highlighting edges and adding contrast, and also increasing noise. In other words, a "smoothing convolution kernel" of a "smoothing function" determines the value of a pixel by examining the neighboring pixels and making the pixel more like its neighbors, which decreases detail. On the other hand, a "sharpening convolution kernel" of a "sharpening function" determines the value of a pixel by examining the neighboring pixels and making the pixel less like its neighbors, which increases detail. Thus, a "smoothing filter" operates to decrease the detail (or focus) of an image, while a "sharpening filter" operates to increase the detail (or focus) of an image.

The "low pass filters" that all implement "gaussian filters" in the image processing system of Lin are all "smoothing" filters. In particular, a "gaussian" filter performs "gaussian smoothing" using a convolution kernel that represents the shape of a gaussian function (i.e., a bell-curve). The gaussian function includes a variable (i.e., the standard deviation) whose value controls the shape of the function (i.e., the height and width of the bell-curve). Thus, the value chosen for this variable determines the degree to which the image will be smoothed (by defining specific values in the corresponding gaussian smoothing convolution kernel).

In Lin, the background portion of the image is scaled using a gaussian filter with kernel coefficients of: "1/10, 1/5, 2/5, 1/5, 1/10", while the text portions of the image are scaled using a gaussian filter with kernel coefficients of: "1/15, 3/15, 7/15, 3/15, 1/15". While the kernel coefficients used in Lin to scale the background portion of the image are different than the kernel coefficients used to scale the text portions of the image, both of the convolution kernels used by Lin have coefficients derived from a gaussian smoothing function and operate to smooth or blur the image portion to which they are applied. In other words, both convolution kernels make the target pixel more like its neighbors, which decreases detail (i.e., smooths the image).

Further, while Lin also uses sigmoid filters on the black and white text image portions, these sigmoid filters do not scale the image. Each sigmoid filter just further filters the image

after it has been scaled by the corresponding low pass filter. Further, a sigmoid filter does not use a convolution kernel at all. A sigmoid filter merely performs a "soft thresholding function" on each pixel individually based on the value of that pixel after scaling to reduce the number of intensity levels displayed in the image. At least, the sigmoid filters are also "smoothing filters" because they operate to decrease the detail in the image (i.e., they raise all levels that are far above the soft threshold to the maximum intensity and lower all levels that are far below the soft threshold to the minimum intensity).

The Examiner has taken the position that "any kernel that sharpens relative to a kernel that smoothes [sic] is a sharpening kernel even though it does not perfectly sharpen." Based on this, the Examiner states that because one of the kernels used in Lin "brings more of the distant pixel" into the target pixel while the other kernel "brings less of the distant pixel" and "more of the middle pixel" into the target pixel, one of the kernels is a "smoothing kernel" and the other kernel is a "sharpening kernel". Such an interpretation of the recited terms "smoothing kernel" and "sharpening kernel" is improper because it completely ignores the established meanings of these terms in the art, or at the least, how these terms are used in the present application.

"Smoothing" and "sharpening" are well defined image effects with specific meanings in both Appellant's specification and in the art of image processing. As explained above, a "smoothing kernel" derived from a "smoothing function" produces a smoothed image by decreasing the detail (or focus) of an image, while a "sharpening kernel" derived from a "sharpening function" produces a sharpened image by increasing the detail (or focus) of an image. The "smoothing kernel" makes the pixel more like its neighbors, which decreases detail and noise, while the "sharpening kernel" makes the pixel less like its neighbors, which increases detail and noise. Just because two convolution kernels have different coefficients does not mean that one of the kernels "smooths the image" while the other of the kernels "sharpens the image". The determination of whether a convolution kernel "smooths" or "sharpens" is made independently for each kernel based on the function from which the kernel coefficients are derived and the effect that the kernel has on an image.

Two kernels with different coefficients are both "smoothing kernels" if each of the kernels is derived from a smoothing function and operates to smooth the image (i.e., by decreasing detail and noise). Conversely, both of the kernels are "sharpening kernels" if each of the kernels is derived from a sharpening function and operates to sharpen the image (i.e., by increasing detail and noise). The smoothing function or sharpening function from which a convolution kernel is derived generally includes a variable that controls the shape of the function, and thus the exact values of the coefficients of the convolution kernel. While the exact values of the coefficients of a "smoothing kernel" affect the degree to which the image is smoothed, given two convolution kernels that are both derived from smoothing functions and that both operate to smooth an image (i.e., remove detail and noise), it is improper to identify one as a "smoothing kernel" and the other as a "sharpening kernel" just because one kernel smooths the image to a greater degree than the other smooths the image (i.e., one operates to remove more detail and noise than the other).

This situation is analogous to the identification of "positive" and "negative" integers. A "positive integer" is defined as an integer that is greater than zero, while a "negative integer" is defined as an integer that is less than zero. The determination of whether an integer is "positive" or "negative" is made independently for each integer based on the value of the integer relative to zero. Two integers with different values are both "negative integers" if each of the integers is less than zero. While the exact values of the integers tells how negative or how positive each is relative to zero, given two integers that are both less than zero, it is improper to identify one as a "negative integer" and the other as a "positive integer" just because one integer is more negative than the other. For example, given the two integers -9 and -3, both are "negative integers" even though -9 is much more negative than -3. It would simply be improper identify -3 as a "positive integer" because it is more positive relative to the value of -9, even though it is not "perfectly positive".

In the same manner, it is improper for the Examiner to identify any kernel that sharpens relative to a kernel that smooths as a sharpening kernel "even though it does not perfectly sharpen," and to use this reasoning to conclude that because one kernel used in Lin "brings more of the distant pixel" into the target pixel while the other kernel "brings less of the distant pixel"

and "more of the middle pixel" into the target pixel, one of the kernels used in Lin is a "smoothing kernel" and the other kernel is a "sharpening kernel". Both of the convolution kernels used in Lin have coefficients derived from a gaussian smoothing function and operate to smooth or blur the image portion to which they are applied. That is, both make the target pixel more like its neighbors, which decreases detail and noise (i.e., smooths the image).

Neither of the convolution kernels used in Lin has coefficients derived from a sharpening function or operates to sharpen an image (i.e., increase detail and noise). While the two sets of coefficients used in Lin are different so as to smooth the image to different degrees, each is clearly representative of a smoothing function and operates to smooth the image and it is improper to define one as a "smoothing kernel" and the other as a "sharpening kernel". Thus, Lin teaches scaling all portions of the source image using low pass filters that use "smoothing convolution kernels" to smooth the image to some degree, and does not teach or suggest using a smoothing kernel and a sharpening kernel on different portions of an image.

In contrast, in the embodiments of the present invention recited in claims 5, 8, 14, 16, and 20, at least one smoothing kernel and at least one sharpening kernel are used in producing a scaled image. More specifically, a local context metric is calculated from a local portion of a source image. A convolution kernel is generated from a plurality of available convolution kernels based on the local context metric, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel. The generated convolution kernel is used to generate at least one pixel of the scaled destination image.

For example, in one exemplary embodiment, the convolution kernels that are used to generate pixels of the scaled image include a "gaussian convolution kernel" derived from a gaussian smoothing function that operates to smooth the image (i.e., decrease detail and noise), and a "cubic convolution kernel" derived from a cubic sharpening function that operates to sharpen the image (i.e., increase detail and noise). See specification at 7:13-8:22. The cubic convolution kernel "produces a sharpened image due to the presence of negative side lobe values". See specification at 2:21-22. These negative side lobe values operate to make the target pixel less like the neighboring pixels so as to increase detail (or focus). In contrast, the gaussian

kernel has all positive side lobe values (like both gaussian kernels of Lin), which operate to make the target pixel more like the neighboring pixels so as to decrease detail (or focus). Thus, the gaussian convolution kernel is a "smoothing kernel" and the cubic convolution kernel is a "sharpening kernel". Accordingly, the image is scaled using a plurality of available convolution kernels that include at least one smoothing kernel and at least one sharpening kernel.

Furthermore, the claimed features of the present invention are not realized even if the teachings of Miyake are incorporated into Lin. Miyake does not teach or suggest the claimed features of the present invention that are absent from Lin. Miyake discloses an image data resolution conversion system, but does not disclose scaling a source image using a plurality of convolution kernels that include at least one smoothing kernel and at least one sharpening kernel.

Both Lin and Miyake, alone or in combination, fail to teach or suggest a method for scaling a source image to produce a scaled destination image in which a convolution kernel is generated from a plurality of available convolution kernels based on the local context metric, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

B. CLAIMS 14, 16, 18, 20, 22, AND 24 ARE PATENTABLE OVER LIN AND
CLAIMS 15 AND 21 ARE PATENTABLE OVER LIN IN VIEW OF MIYAKE

Appellant respectfully submits that claims 14, 16, 18, 20, 22, and 24 (of Group II) are patentable over Lin and claims 15 and 21 (of Group II) are patentable over Lin in view of Miyake because neither Lin nor Miyake teaches or suggests an image scaling device for producing a scaled destination image that includes a kernel generator for generating a current convolution kernel from a plurality of available convolution kernels based on a local context metric, and a scaler for using coefficients of the current convolution kernel to generate at least one pixel of the scaled destination image, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

Claim 14 recites an image scaling device for receiving pixels of a source image and outputting pixels of a scaled destination image that includes a kernel generator for generating a current convolution kernel from a plurality of available convolution kernels based on a local context metric calculated by a context sensor, and a scaler for using coefficients of the current convolution kernel to generate at least one pixel of the scaled destination image, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

As discussed above, both Lin and Miyake, alone or in combination, fail to teach or suggest a method for scaling a source image to produce a scaled destination image in which a convolution kernel is generated from a plurality of available convolution kernels based on the local context metric, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

Further, even if Lin and Miyake, alone or in combination, were proper to somehow generally a method for scaling a source image to produce a scaled destination image in which at least one smoothing kernel and at least one sharpening kernel are used, both Lin and Miyake fail to teach an image scaling device that includes the specific kernel generator and scalar that are recited in claim 14.

In the embodiment of the present invention recited in claim 14, the image scaling device that includes a context sensor, a kernel generator, and a scalar. The context sensor calculates a local context metric based on local source image pixels, and the kernel generator generates a current convolution kernel from a plurality of available convolution kernels based on the local context metric calculated by the context sensor. The scaler receives the coefficients of the current convolution kernel from the kernel generator, and uses the coefficients to generate at least one pixel of the scaled destination image from pixels of the source image. The available convolution kernels include at least one smoothing kernel and at least one sharpening kernel.

Thus, the image scaling device includes a kernel generator that generates a current convolution kernel from a plurality of available convolution kernels that include at least one smoothing kernel and at least one sharpening kernel. Further, the scalar receives from the kernel

generator the coefficients of the convolution kernel that was generated by the kernel generator, and uses these received coefficients to scale the image. Neither Lin nor Miyake teaches or suggests the recited kernel generator and scaler for use in an image scaling device that employs multiple convolution kernels.

Both Lin and Miyake, alone or in combination, fail to teach or suggest an image scaling device that includes a kernel generator for generating a current convolution kernel from a plurality of available convolution kernels based on a local context metric, and a scaler for receiving coefficients of the current convolution kernel from the kernel generator and using these coefficients to generate at least one pixel of the scaled destination image, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

C. CLAIMS 7, 13, 19, AND 25 ARE PATENTABLE OVER LIN IN VIEW OF MIYAKE

Appellant respectfully submits that claims 7, 13, 19, and 25 (of Group II) are patentable over Lin in view of Miyake because neither Lin nor Miyake teaches or suggests a method or device for producing a scaled destination image in which the available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

Claim 7 recites a method for scaling a source image to produce a scaled destination image in which a convolution kernel is generated from a plurality of available convolution kernels based on a calculated local context metric, and the generated convolution kernel is used to generate at least one pixel of the scaled destination image, and the available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

As discussed above, both Lin and Miyake, alone or in combination, fail to teach or suggest a method for scaling a source image to produce a scaled destination image in which a convolution kernel is generated from a plurality of available convolution kernels based on the local context metric, with the available convolution kernels including at least one smoothing kernel and at least one sharpening kernel.

Further, even if Lin and Miyake, alone or in combination, were proper to somehow generally a method for scaling a source image to produce a scaled destination image in which at least one smoothing kernel and at least one sharpening kernel are used, both Lin and Miyake fail to teach an image scaling method or device in which there are used a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel. Miyake only teaches using smoothing filters in scaling the image.

In contrast, in the embodiment of the present invention recited in claim 7, a local context metric is calculated from a local portion of a source image. A convolution kernel is generated from a plurality of available convolution kernels based on the local context metric, and the generated convolution kernel is used to generate at least one pixel of the scaled destination image. The available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

Thus, in the scaling method, the image is scaled using a plurality of available convolution kernels that include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel. Neither Lin nor Miyake teaches or suggests this recited feature.

Both Lin and Miyake, alone or in combination, fail to teach or suggest a method or device for producing a scaled destination image in which the available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

In view of the foregoing, it is respectfully submitted that the application and the claims are in condition for allowance. Reversal of the final rejection of claims 62-108 is respectfully requested.

Respectfully submitted,

Date: August 9, 2004

By:



Stephen Bongini
Registration No. 40,917
Attorney for Appellant

FLEIT, KAIN, GIBBONS,
GUTMAN, BONGINI & BIANCO P.L.
One Boca Commerce Center
551 Northwest 77th Street, Suite 111
Boca Raton, Florida 33487
Telephone: (561) 989-9811
Facsimile: (561) 989-9812

9. APPENDIX

3. The method as defined in claim 5, further comprising the step of:
storing all available convolution kernels in a memory,
wherein in the generating step, one of the stored convolution kernels is selected based on
the calculated local context metric.
5. A method for scaling a source image to produce a scaled destination image, said method
comprising the steps of:
calculating a local context metric from a local portion of the source image;
generating a convolution kernel from a plurality of available convolution kernels based on
the calculated local context metric; and
using the generated convolution kernel to generate at least one pixel of the scaled
destination image, the scaled destination image having a different resolution than the source
image,
wherein the available convolution kernels include at least one smoothing kernel and at
least one sharpening kernel.
6. The method as defined in claim 5, wherein the local context metric has more than two
possible values.

7. The method as defined in claim 6, wherein the available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

8. A machine-readable medium encoded with a program for scaling a source image to produce a scaled destination image, said program containing instructions for performing the steps of:

calculating a local context metric from a local portion of the source image;
generating a convolution kernel from a plurality of available convolution kernels based on the calculated local context metric; and
using the generated convolution kernel to generate at least one pixel of the scaled destination image, the scaled destination image having a different resolution than the source image,

wherein the available convolution kernels include at least one smoothing kernel and at least one sharpening kernel.

9. The machine-readable medium as defined in claim 8, wherein said program further contains instructions for performing the step of repeating the calculating, generating, and using steps for each pixel in the destination image.

10. The machine-readable medium as defined in claim 8, wherein said program further contains instructions for performing the step of:

storing all available convolution kernels in a memory,

wherein in the generating step, one of the stored convolution kernels is selected based on the calculated local context metric.

12. The machine-readable medium as defined in claim 8, wherein the local context metric has more than two possible values.

13. The machine-readable medium as defined in claim 12, wherein the available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

14. An image scaling device that receives pixels of a source image and outputs pixels of a scaled destination image, said image scaling device comprising:

a context sensor for calculating a local context metric based on local source image pixels;
a kernel generator coupled to the context sensor, the kernel generator generating a current convolution kernel from a plurality of available convolution kernels based on the local context metric calculated by the context sensor; and
a scaler coupled to the kernel generator, the scaler receiving the coefficients of the current convolution kernel from the kernel generator, and using the coefficients to generate at least one pixel of the scaled destination image from pixels of the source image, the scaled destination image having a different resolution than the source image,

wherein the available convolution kernels include at least one smoothing kernel and at least one sharpening kernel.

15. The image scaling device as defined in claim 14, wherein the context sensor calculates a local context metric for each pixel in the destination image.

16. An image scaling device that receives pixels of a source image and outputs pixels of a scaled destination image, said image scaling device comprising:

a context sensor for calculating a local context metric based on local source image pixels;

a kernel generator coupled to the context sensor, the kernel generator generating a current convolution kernel from a plurality of available convolution kernels based on the local context metric calculated by the context sensor; and

a scaler coupled to the kernel generator, the scaler receiving the coefficients of the current convolution kernel from the kernel generator, and using the coefficients to generate at least one pixel of the scaled destination image from pixels of the source image, the scaled destination image having a different resolution than the source image,

wherein the available convolution kernels include at least one smoothing kernel and at least one sharpening kernel,

the kernel generator stores all available convolution kernels, and

the kernel generator selects one of the stored convolution kernels as the current convolution kernel based on the calculated local context metric.

18. The image scaling device as defined in claim 14, wherein the local context metric has more than two possible values.

19. The image scaling device as defined in claim 18, wherein the available convolution kernels include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other kernels that provide a transition between the complete sharpening kernel and the complete smoothing kernel.

20. A display device that receives source image pixels and displays a scaled destination image, said display device comprising:

a context sensor for calculating a local context metric based on local source image pixels;
a kernel generator coupled to the context sensor, the kernel generator generating a current convolution kernel from a plurality of available convolution kernels based on the local context metric calculated by the context sensor;

a scaler coupled to the kernel generator, the scaler receiving the coefficients of the current convolution kernel from the kernel generator, the scaler using the coefficients to generate at least one pixel of the scaled destination image from pixels of the source image, the scaled destination image having a different resolution than the source image; and

a display for displaying the scaled destination image,
wherein the available convolution kernels include at least one smoothing kernel and at least one sharpening kernel.

21. The display device as defined in claim 20, wherein the context sensor calculates a local context metric for each pixel in the destination image.

22. The display device as defined in claim 20,
wherein the kernel generator stores all available convolution kernels, and
the kernel generator selects one of the stored convolution kernels as the current
convolution kernel based on the calculated local context metric.
24. The display device as defined in claim 20, wherein the display is an LCD display.
25. The display device as defined in claim 20, wherein the available convolution kernels
include a complete smoothing kernel, a complete sharpening kernel, and a plurality of other
kernels that provide a transition between the complete sharpening kernel and the complete
smoothing kernel.